Effect of Incomplete Meta-dataset on Average Ranking Method

Salisu Mamman Abdulrahman SALISU.ABDUL@GMAIL.COM
LIAAD - INESC TEC/Faculdade de Ciências da Universidade do Porto

Pavel Brazdil PBRAZDIL@INESCPORTO.PT
LIAAD - INESC TEC/Faculdade de Economia, Universidade do Porto

arXiv:1608.06845v1 [cs.AI] 24 Aug 2016

Abstract

One of the simplest metalearning methods is the average ranking method. This method uses metadata in the form of test results of a given set of algorithms on given set of datasets and calculates an average rank for each algorithm. The ranks are used to construct the average ranking. We investigate the problem of how the process of generating the average ranking is affected by incomplete metadata including fewer test results. This issue is relevant, because if we could show that incomplete metadata does not affect the final results much, we could explore it in future design. We could simply conduct fewer tests and save thus computation time. In this paper[1] we describe an upgraded average ranking method that is capable of dealing with incomplete metadata. Our results show that the proposed method is relatively robust to omission in test results in the meta datasets.

Keywords: Average Ranking, Aggregation of Rankings, Incomplete Metadata

1. Introduction

A large number of data mining algorithms exist, rooted in the fields of machine learning, statistics, pattern recognition, artificial intelligence. The task to recommend the most suitable algorithms has thus become rather challenging. The algorithm selection problem, originally described by Rice (1976), has attracted a great deal of attention, as it endeavours to select and apply the best or near best algorithm(s) for a given task (Brazdil et al. (2008); Smith-Miles (2008)). We address the problem of robustness of one particular version of an average ranking method that uses incomplete rankings as input. These arise if we have incomplete test results in the meta-dataset. We have investigated how much the performance degrades under such circumstances.

The remainder of this paper is organized as follows. In the next section we present an overview of existing work in related areas. Section 3 provide details about the proposed aggregation method for incomplete meta dataset, the experimental results and future work.

1. This paper is a slightly updated version of the paper presented at AutoML workshop at ICML 2016, New York

© 2016 S.M. Abdulrahman & P. Brazdil.

2. Related Work

In this paper we are addressing a particular case of the algorithm selection problem, oriented towards the selection of classification algorithms. Various researchers addressed this problem in the course of the last 25 years. One approach to algorithm selection/recommendation relies on metalearning. The simplest method uses just performance results on different datasets in the form of rankings. The rankings are then aggregated to obtain a single aggregated ranking. So the aggregated ranking can be used as a simple model that can be followed by the user to test the top candidates to identify the algorithm to be used. This strategy is sometimes referred to as the *Top-N* strategy (Brazdil et al. (2008)).

A more advanced approach often considered as the *classical metalearning approach* uses, in addition to performance results, also a set of measures that characterize datasets (Pfahringer et al. (2000); Brazdil et al. (2008); Smith-Miles (2008)). However, this line is not followed up here.

Aggregation of rankings involving *complete rankings* is a simple matter. Normally it just involves calculating the average rank for all items in the ranking (Lin, 2010). Complete rankings are those in which k items are ranked N times and no value in this set is missing. Incomplete rankings arise when only some ranks are known in some of the rankings. These arise quite often in practice. Many diverse methods exist for aggregation of incomplete rankings. According to Lin (2010) methods applicable to long lists can be divided into three categories: Heuristic algorithms, Markov chain methods and stochastic optimization methods. The last category includes, for instance, *Cross Entropy Monte Carlo, CEMC* method.

Some of the approaches require that the elements that do not appear in list L_i of k elements be attributed a concrete rank (e.g. $k+1$). This does not seem to be correct. We should not be forced to assume that some information exists, if in fact we have none. We have considered using a package of R *RankAggreg* (Pihur et al. (2014)), but unfortunately we would have to attribute a concrete rank (e.g. $k+1$) to all missing elements. We have therefore developed a simple heuristic method based on Borda's method described in Lin (2010). In our view it serves well our purpose. As Lin (2010) pointed out simple methods often compete quite well with other more complex approaches.

3. Effect of Incomplete Meta-dataset on Average Ranking Method

Our aim is to investigate the issue of how the generation of the average ranking is affected by incomplete test results in the meta-dataset available. Here we focus on rankings obtained on the basis of accuracies. We wish to see how robust the method is to omissions in the meta-dataset. This issue is relevant because the meta-dataset that has been gathered by researchers is very often incomplete. Here we consider two different ways in which the meta-dataset can be incomplete: First, the test results on some datasets may be completely missing. Second, there may be certain proportion of omissions in the test results of some algorithms on each dataset.

The expectation is that the performance of the average ranking method would degrade when less information is available. However, an interesting question is how grave the degradation is. The answer to this issue is not straightforward, as it depends greatly on how diverse the datasets are and how this affects the rankings of algorithms. If the rankings are

very similar, then we expect that the omissions would not make much difference. So the issue of the effects of omissions needs to be relativized. To do this we will investigate the following issues:

- Effects of missing test results on X% of datasets (alternative MTD);
- Effects of missing X% of test results of algorithms on each dataset (alternative MTA).

If the performance drop of alternative MTA were not too different from the drop of alternative MTD, than we could conclude that X% of omissions is not unduly degrading the performance and hence the method of average ranking is relatively robust. Each of these alternatives is discussed in more detail below.

Missing all test results on some datasets (alternative MTD): This strategy involves randomly omitting all test results on a given proportion of datasets from our meta-dataset. An example of this scenario is depicted in Table 1. In this example the test results on datasets D_2 and D_5 are completely missing. The aim is to show how much the average ranking degrades due these missing results.

Table 1: Missing test results on certain percentage of datasets (MTD)

Algorithms	D_1	D_2	D_3	D_4	D_5	D_6
a_1	0.85		0.77	0.98		0.82
a_2	0.95		0.67	0.68		0.72
a_3	0.63		0.55	0.89		0.46
a_4	0.45		0.34	0.58		0.63
a_5	0.78		0.61	0.34		0.97
a_6	0.67		0.70	0.89		0.22

Missing some algorithm test results on each dataset (alternative MTA): Here the aim is to drop a certain proportion of test results on each dataset. The omissions are simply distributed uniformly across all datasets. That is, the probability that the test result of algorithm a_i is missing is the same irrespective of which algorithm is chosen. An example of this scenario is depicted in Table 2. The proportion of datasets is a parameter of the

Table 2: Missing test results on certain percentage of algorithms (MTA)

Algorithms	D_1	D_2	D_3	D_4	D_5	D_6
a_1	0.85	0.77		0.98		0.82
a_2		055	0.67	0.68	0.66	
a_3	0.63		0.55	0.89		0.46
a_4	0.45	0.52	0.34		0.44	0.63
a_5	0.78	0.87	0.61	0.34	0.42	
a_6		0.99		0.89		0.22

method. Here we use the values shown in Table 3. The meta-dataset used in this study is

described further on in Section 3.2. This dataset was used to obtain a new one in which some datasets would be chosen at random and all test results obliterated. The resulting dataset was used to construct the average ranking. Each ranking was then used to construct a *loss-time curve* described further on (in Section 3.2). The whole process was repeated 10 times. This way we would obtain 10 loss-time curves, which would be aggregated into a single loss-time curve. Our aim is to upgrade the average ranking method to be able to deal with incomplete rankings. The enhanced method is described in the next section.

Table 3: Percentages of omissions and the numbers of datasets and algorithms used

Omissions %	0	5	10	20	50	90	95
No of datasets used	38	36	34	30	19	4	2
No of tests per dataset	53	50	48	43	26	5	3

3.1. Aggregation Method for Incomplete Rankings (AR-MTA)

Before describing the method for the calculation of the average ranking that can process incomplete rankings, let us consider a motivating example (see Table 4), illustrating why we cannot simply use the usual average ranking method (Lin (2010)), often used in comparative studies in machine learning literature. Let us compare the rankings R_1 (Table 4a) and R_2

Table 4: An example of two rankings R_1 and R_2 and the aggregated ranking R^A

(a)

R_1	Rank
a_1	1
a_3	2
a_4	3
a_2	4
a_6	5
a_5	6

(b)

R_2	Rank
a_2	1
a_1	2

(c)

R^A	Rank	Weight
a_1	1.67	1.2
a_3	2	1
a_4	3	1
a_2	3.5	1.2
a_6	5	1
a_5	6	1

(Table 4b). We note that algorithm a_2 is ranked 4 in ranking R_1, but has rank 1 in ranking R_2. If we used the usual method, the final ranking of a_2 would be the mean of the two ranks, i.e. (4+1)/2=2.5. This seems intuitively not right, as the information in ranking R_2 is incomplete. If we carry out just one test and obtain ranking R_2 as a result, this information is obviously inferior to having conducted more tests leading to ranking R_1. So these observations suggest that the number of tests should be taken into account to set a weight to characterize the individual elements of the ranking.

In our method the weight is calculated using the expression $(N-1)/(Nmax-1)$, where N represents the number of filled-in elements in the ranking and $Nmax$ the maximum number of elements that could be filled-in. The number $N-1$ represents the number of non-transitive relations between any element in the ranking to any other element in the ranking. So, for instance, in ranking R_1 $N=6$ and $Nmax=6$. Therefore, the weight of

each element in the ranking is $5/5 = 1$. We note that $N-1$ (i.e. 5), represents the number of non-transitive relations in the ranking, namely $a_1 > a_3$, $a_3 > a_4$, .. , $a_6 > a_5$. Here $a_i > a_j$ is used to indicate that a_i is preferred to a_j.

Let us consider the incomplete ranking R_2. Suppose we know a priori that the ranking could include 6 elements and so $Nmax = 6$, as in the previous case. Then the weight of each element will be $(N-1)/(Nmax-1) = 1/5 = 0.2$. The notion of *weight* captures the fact that ranking R_2 provides less information than ranking R_1. We need this concept in the process of calculating the average ranking.

Our upgraded version of the aggregation method for incomplete rankings involves the initialization step, which consists of reading-in the first ranking and initializing the average ranking R^A. Then in each subsequent step a new ranking is read-in and aggregated with the average ranking, producing a new average ranking. The aggregation is done by going through all elements in the ranking, one by one. If the element appears in both the aggregated ranking and the read-in ranking, its rank is recalculated as a weighted average of the two ranks:

$$r_i^A := r_i^A * w_i^A/(w_i^A + w_i^j) + r_i^j * w_i^j/(w_i^A + w_i^j) \quad (1)$$

where r_i^A represents the rank of element i in the aggregated ranking and r_i^j the rank of the element in the ranking j that is being processed and w_i^A and w_i^j represent the corresponding weights. The weight is updated as follows: $w_i^A := (w_i^A + w_i^j)$. If the element appears in the aggregated ranking, but not in the new read-in ranking, the new aggregated ranking is made equal to the previous one. The weight of the element is also kept unchanged.

Suppose the aim is to aggregate ranking R_1 and R_2 shown before. The new rank of a_2 will be $r_2^A = 4 * 1/1.2 + 1*0.2/1.2 = 3.5$. The weight will be $w_2^A = 1 + 0.2 = 1.2$. The final aggregated ranking of rankings R_1 and R_2 is shown in Table 4c.

3.2. Results on the Effects of Omissions in the Meta-Dataset

The data used in the experiments involves the meta-dataset constructed from evaluation results retrieved from OpenML (Vanschoren et al. (2014)), a collaborative science platform for machine learning. This dataset contains the results of 53 parameterized classification algorithms from the Weka workbench (Hall et al. (2009)) on 39 datasets[2]. In the leave-one-out mode, 38 datasets are used to generate the model (e.g. average ranking), while the dataset left out is used for evaluation.

Characterization of our Meta-Dataset: We are interested to analyze rankings of classification algorithms on different datasets and in particular how these differ for pairs of datasets, using Spearman correlation coefficient. Fig.1 shows a histogram characterizing the meta-dataset used. The histogram is accompanied by *expected value*, *standard deviation* and *coefficient of variation* calculated as the ratio of standard deviation to the expected value (mean) (Witten and Frank (2005)). These measures are shown in Table 5.

Results: The aim is to investigate how certain omissions in the meta-datasets affect the performance. The results are presented in the form of loss-time curve (van Rijn et al. (2015)) which show how *performance loss* depends on time. The *loss* is calculated as the difference between the performance of the best algorithm identified so far when following

2. Full details: http://www.openml.org/project/tag/ActiveTestingSamples/u/1

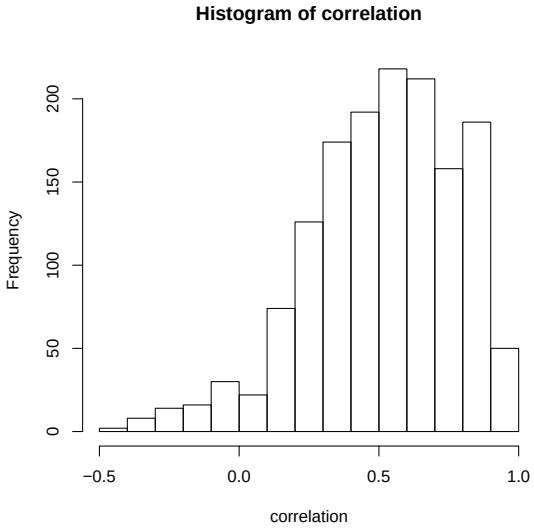

Figure 1: Spearman correlation between rankings of pairs of datasets.

Table 5: Measures characterizing the histogram of correlations in Fig. 1

Measure %	Expected Value	Standard Deviation	Coefficient of Variation
Value	0.5134	0.2663	51.86%

the ranking to the ideal choice. Each loss-time curve can be characterized by a number representing the *mean loss* in a given interval. We want this *mean interval loss* (MIL) to be as low as possible. This characteristic is similar to AUC, but there is an important difference. When talking about AUCs, the x-axis values spans between 0 and 1. Our loss-time curves span between some T_{min} and T_{max} and both values depend on the user. Typically the user searching for a suitable algorithm would not worry about quite short times. In the experiments here we have set T_{min} to 10 seconds. The value of T_{max} was set to 10^4 seconds, i.e. about 2.78 hours.

Table 6 presents the results for both alternative *AR-MTD* and *AR-MTA* in terms of mean interval loss (MIL). The values for ordinary average ranking method, *AR*, are also shown, as this method serves as a baseline. Fig.2 shows the loss-time curves for the alternatives *AR-MTD* and *AR-MTA* when the number of omissions is 90%. The values of *AR* relative to 90% omissions are also shown for comparison. Not all loss-time curves are shown, as the figure would be rather cluttered.

Our results show that the alternative *AR-MTA* achieves far better results than the baseline method *AR*. We note also that although the proposed method *AR-MTA* is worse than the *AR-MTD* counterpart, the difference is negligible for many different values of omissions in the range from 5% till 50%. Only when we get to rather extreme values such as 90%, the difference is noticeable. But even in this case the curve of *AR-MTA* follows the curve of

Table 6: Mean interval loss (MIL) values for different percentage of omissions

Method \ Omission%	0%	5%	10%	20%	50%	90%	95%
AR	1.213	2.306	2.435	2.785	2.507	2.302	2.360
AR-MTD	1.213	1.129	0.812	0.937	1.046	0.981	1.387
AR-MTA	1.213	0.881	0.971	0.972	1.031	1.473	1.799

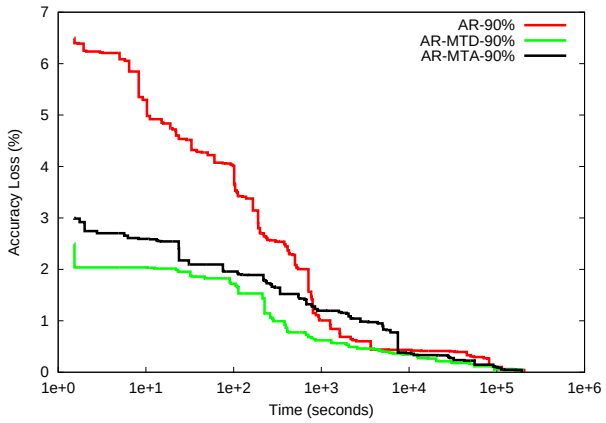

Figure 2: Comparison of *AR-MTA* with *AR-MTD* and the baseline method *AR* for 90% of omissions.

AR-MTD much more closely than *AR*. These results indicate that the proposed average ranking method is relatively robust to omissions.

Future work: As the incomplete meta-dataset does not affect much the final ranking and the corresponding loss, this could be explored in future design of experiments, when gathering the test results. We could investigate approaches that permit to consider also the costs (time) of off-line tests. Their cost (time) could be set to some fraction of the cost of on-line test (i.e. tests on a new dataset), but not really ignored altogether.

Acknowledgements
This work is supported by TETFund 2012 Intervention for Kano University of Science and Technology, Wudil, Kano State, Nigeria for PhD Overseas Training from the Federal Government of Nigeria.

The authors also acknowledge the support of project *NanoSTIMA: Macro-to-Nano Human Sensing: Towards Integrated Multimodal Health Monitoring and Analytics/NORTE-01-0145-FEDER-000016*, which is financed by the North Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 Partnership Agreement, and through the European Regional Development Fund (ERDF).

References

Pavel Brazdil, Christophe Giraud-Carrier, Carlos Soares, and Ricardo Vilalta. *Metalearning: Applications to data mining*. Springer Science & Business Media, 2008.

M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.

S. Lin. Rank aggregation methods. *WIREs Computational Statistics*, 2:555–570, 2010.

B. Pfahringer, H. Bensusan, and Ch. Giraud-Carrier. Tell me who can learn you and I can tell you who you are: Landmarking various learning algorithms. In *Proc. of the 17th Int. Conf. on Machine Learning*, pages 743–750, 2000.

V. Pihur, S. Datta, and S. Susmita Datta. RankAggreg, an R package for weighted rank aggregation. Department of Bioinformatics and Biostatistics, University of Louisville, http://vpihur.com/biostat, 2014.

J. R. Rice. The Algorithm Selection Problem. *Advances in Computers*, 15:65–118, 1976.

K. A. Smith-Miles. Cross-disciplinary Perspectives on Meta-Learning for Algorithm Selection. *ACM Computing Surveys (CSUR)*, 41(1):6:1–6:25, 2008.

J. N. van Rijn, S. M. Abdulrahman, P. Brazdil, and J. Vanschoren. Fast Algorithm Selection using Learning Curves. In *Advances in Intelligent Data Analysis XIV*. Springer, 2015.

J. Vanschoren, J. N. van Rijn, B. Bischl, and L. Torgo. OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014.

Ian H. Witten and Eibe Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2005.